Device Class Power Management Reference Specification

Storage Device Class

Draft proposal
v0.9

Table of Contents

Scope 3
General Device Power Management Considerations 3
Hard Disk and CD-ROM Device Power State Definitions 4
Floppy Disk Device Power State Definitions 4
Hard Disk, Floppy Disk and CD-ROM Device Power Conservation Policy 6
IDE Channel Power State Definitions 6
IDE Channel Power Conservation Policy 6
Storage Device Wake-up Events 6
Minimum Storage Device Power Capabilities ..... **Error! Bookmark not defined.**

Revision History

Revision	Date	Comments
0.0	3/25/96	Initial proposal for consideration
0.9	9/20/96	Complete state definitions

Scope

This specification defines the behavior of Storage devices as it relates to power management, and, specifically, to the four device power states defined for the OnNow Architecture. This specification specifically covers ATA Hard Disks, Floppy disks, ATAPI and SCSI CD-ROMs, and the IDE channel. It is intended that storage vendors and system makers will be able to design consistent power-manageable products, and that OS vendors will be able to implement an appropriate storage power management policy based on the contents of this specification.

Abstract

Until now, hardware vendors have been attacking power management from the perspective of saving as much power as possible when the device changes from its full-power state. We are attacking the problem from a different perspective - there should be a power management state that saves as much power as possible *without causing noticeable effects for the user*. This will give the user both power savings and good performance. There will be lower power modes beyond this that are fairly analogous to existing power modes today.

General Device Power Management Considerations

In the OnNow architecture, power management of individual devices is the responsibility of a policy owner in the Operating System, generally a class-specific driver. This policy-owner will implement a power conservation policy that is appropriate for devices in its class. The policy will operate in conjunction with a global system power policy implemented in the operating system (i.e. is the system Working or Sleeping?). In general, the device-class power conservation policy strives to reduce power consumption while the system is Working by transitioning amongst various available power states according to device usage. Since the policy-owner in the Operating System has very specific knowledge of when a device is in use, or potentially in use, there is no need for hardware timers or such to determine when to make these transitions. Similarly, this level of understanding of device usage makes it possible to use fewer device power states. Generally, intermediate states attempt to draw a compromise between latency and consumption due to the uncertainty of actual device usage. With the increased knowledge in the OS, crisp decisions can be made about whether the device is needed at all. With this ability to turn devices off more frequently, the benefit of having intermediate states diminishes.

The policy-owner also determines what class-specific events can cause the system to transition from Sleeping to Working, and enables this functionality based on application or user requests. Note that the definition of the wake-up events that each class supports will influence the system's global power policy in terms of the level of power conservation the Sleeping state can attain while still meeting wake-up latency requirements set by applications or the user.

In the OnNow architecture, bus drivers also implement power policy for their bus class (e.g. PCI, USB, etc..). In general, the Bus driver has responsibility for tracking the device power states of all devices on its bus, and transitioning the Bus itself to only those power states that are consistent with those of its devices. This means that the Bus state can be no lower than the highest state of one of it's devices. However, enabled wake-up events can affect this as well. For example if a particular device is in the D2 state and set to wake-up the system, and the bus can only forward wake-up requests while in the D1 state, then the Bus must remain in the D1 state even if all devices are in a lower state.

Device power state transitions are explicitly commanded by the driver and invoked through bus-specific mechanisms (e.g. ATA Standby command, USB Suspend, etc..). In some cases, bus-specific mechanisms are not available and device-specific mechanisms must be used. Note that the explicit command for entering the D3 state may be the removal of power.

The following definitions apply to devices of all classes:

- **D0**: Device is on and running. It is receiving full power from the system, and is delivering full functionality to the user.

- **D1**: Class-specific low-power state (defined below) in which device context may or may not be lost. Buses in D1 cannot do anything to the bus which would force devices on that bus to lose context.
- **D2**: Class-specific low-power state (defined below) in which device context may or may not be lost. Attains greater power savings than D1. Buses in D2 may cause devices on that bus to lose some context (e.g. the bus reduces power supplied to the bus). Devices in D2 must be prepared for the bus to be in D2 (or higher).
- **D3**: Device is off and not running. Device context is lost. Power may be removed from the device.

Any device context lost must be restored by the device driver when returning the device to the D0 state.

Hard Disk, CD-ROM and IDE/ATAPI removable storage Device Power State Definitions

D0

- Drive controller (i.e. interface and control electronics): functional
- Interface mode context (i.e. communications timings): programmed
- Drive motor (i.e. spindle): turning
- Laser (if any): On

D1

- Drive controller (i.e. interface and control electronics): functional
- Interface mode context (i.e. communications timings): preserved
- Drive motor (i.e. spindle): stopped, with fast-start mode enabled, if available.
- Laser (if any): Off
  NOTE: For ATA devices, this state is invoked by the Standby Immediate command

Latency to return to D0:	3 seconds (or less)
Power consumption in D1:	No more than 80% of power consumed in D0

D2

This state is not defined for storage devices. Use D3 instead.

D3 (Power may be removed)

- Drive controller (i.e. interface and control electronics): not functional; context lost
- Interface mode (i.e. communications timings): not preserved
- Drive motor (i.e. spindle): stopped
- Laser (if any): Off
  NOTE: For ATA devices, this state is invoked by the Sleep Immediate command

Power consumption in D3:	No more than 10% of power consumed in D0

Floppy Disk Device Power State Definitions

D0

- Drive controller (i.e. interface and control electronics): functional
- Drive motor (i.e. spindle): turning

D1

This state is not defined for storage devices. Use D3 instead

D2

This state is not defined for storage devices. Use D3 instead

D3 (Power may be removed)

- Drive controller (i.e. interface and control electronics): not functional; context lost
- Drive motor (i.e. spindle): stopped

Hard Disk, Floppy Disk, IDE/ATAPI removable storage device and CD-ROM Device Power Conservation Policy

Present State	Next State	Cause
D3	D0	Device usage (High-priority I/O)
D0	D1*	Device inactivity (no high-priority I/O) for a period of time (T1)
D0	D3	• Device inactivity (no high-priority I/O) for a period of time (T2 =>T1) • System enters sleeping state
D1*	D0	Device usage (High-priority I/O)

*If available

NOTE: For ATA, the D3-to-D0 transition requires a reset of the IDE channel. This means that both devices on a channel must be placed into D3 at the same time.

IDE Channel Power State Definitions

D0

- Adapter: Functional
- Adapter interface mode (i.e. communications timings): Programmed
- Power to Bus (all devices connected to it): On

D1

This state is not defined for the IDE Channel.

D2

This state is not defined for the IDE Channel.

D3 (Power may be removed)

- Adapter: Non-functional
- Adapter interface mode (i.e. communications timings): not preserved
- Power to Bus (all devices connected to it): Off

IDE Channel Power Conservation Policy

Present State	Next State	Cause
D3	D0	Any device on the channel needing to transition to a state other than state D3
D0	D3	• All devices on the channel in state D3

SCSI Channel Power State Definitions

<<This section is undefined. Please send comments on possibilities for SCSI power state definitions to power@microsoft.com. Without help from outside vendors, we will not be able to implement OnNow on SCSI.>>

Storage Device Wake-up Events

Storage devices with removable media can, optionally, signal wake-up upon insertion of media using their bus-specific notification mechanism. There are no other wake-up events defined for Storage devices.